

Empirical correction of a toy climate model

Nicholas A. Allgaier, Kameron D. Harris, and Christopher M. Danforth
*Department of Mathematics and Statistics, Vermont Advanced Computing Center,
Complex Systems Center, The University of Vermont*

(Dated: July 15, 2011)

Improving the accuracy of forecast models for physical systems such as the atmosphere is a crucial ongoing effort. Errors in state estimation for these often highly nonlinear systems has been the primary focus of recent research, but as that error has been successfully diminished, the role of model error in forecast uncertainty has duly increased. The present study is an investigation of a particular empirical correction procedure that is of special interest because it considers the model a “black box”, and therefore can be applied widely with little modification. The procedure involves the comparison of short model forecasts with a reference “truth” system during a training period in order to calculate systematic (1) state-independent model bias and (2) state-dependent error patterns. An estimate of the likelihood of the latter error component is computed from the current state at every timestep of model integration. The effectiveness of this technique is explored in two experiments: (1) a perfect model scenario, in which models have the same structure and dynamics as the true system, differing only in parameter values; and (2) a more realistic scenario, in which models are structurally different (in dynamics, dimension, and parameterization) from the target system. In each case, the results suggest that the correction procedure is more effective for reducing error and prolonging forecast usefulness than parameter tuning. However, the cost of this increase in average forecast accuracy is the creation of substantial qualitative differences between the dynamics of the corrected model and the true system. A method to mitigate the structural damage caused by empirical correction and further increase forecast accuracy is presented.

I. INTRODUCTION

Advances in computational power and increasingly accurate techniques for estimating the current state of the Earth’s atmosphere have significantly improved numerical weather prediction (NWP) [1–3]. As state estimation error is reduced due to improved methods of data assimilation, error in the model tendency plays an increasing role in the uncertainty of predictions at every temporal and physical scale [4–6].

In 1978, Leith introduced a statistical technique to correct model tendency error, in which short model forecasts are compared to a time series of reference “truth” to estimate both state-independent model bias, and state-dependent error components which are approximated by a least-squares linear function of the model state [7]. More recently, empirical correction has been employed with success in atmospheric models with relatively few degrees of freedom (e.g. $N = O(10^2)$ in [8]), and low-dimensional modifications of the technique involving, for example, singular value decomposition (SVD) of the state-dependent correction operator (the least-squares linear function proposed by Leith) have proven successful in models with as many as $O(10^5)$ degrees of freedom [9]. In this study, we apply the original technique developed by Leith to simple three-dimensional Lorenz models [10], where in addition to testing the effectiveness of empirical correction, we aim to understand the dynamical ramifications of a statistical approach to the correction of model tendency error.

The model tendency $M(\mathbf{x})$ is defined as the change in state variables over one timestep of numerical integra-

tion, which we denote as a time derivative:

$$M(\mathbf{x}) \equiv \dot{\mathbf{x}}^M \quad (1)$$

where \mathbf{x} is the atmospheric state-vector, typically with $O(10^{10})$ degrees of freedom for NWP. Note that \mathbf{x}^M represents the model state, whereas \mathbf{x}^T will represent the true system state, in terms of model variables.

Given the state \mathbf{x} of a physical system, M consists of all the known physics, forcings, and parameterizations of sub grid-scale processes. To make a one-timestep model forecast, we approximate the true change in state variables over that time by the model tendency, $\dot{\mathbf{x}}^T \approx M(\mathbf{x})$, and the error δM in that approximation is called the *tendency error*:

$$\delta M = \dot{\mathbf{x}}^T - M(\mathbf{x}) \quad (2)$$

Even with perfect estimates of the current state of the atmosphere, the model tendency error would quickly separate forecasts from the truth, due to the atmosphere’s chaotic dynamics [11]. As a result, techniques for reducing the model tendency error represent a current and active research area, and those that are applicable independent of the specific model are of special interest. Clearly, one would like to improve the physics represented by M from first principles. In what follows, we assume this improvement has met the limit of diminishing returns, and move towards a statistical approach.

The general strategy of empirical correction is to compare short forecasts generated by a model to observations of the system being modeled over some training period. If the state-space of the system is well represented by the training period, and the model is a reasonable approximation of the true system, the forecast error statistics

can be used to create an empirical correction that pushes the model closer to the truth *at each timestep* of numerical integration. Adjusting the model every timestep reduces the nonlinear growth of tendency error, providing more effective error reduction than a posteriori statistical correction [12]. This strategy is similar to nudging or Newtonian relaxation in a data assimilation (DA) context, where one is assimilating observations, except that here we are nudging with predicted, rather than observed, forecast error.

The present study is an investigation of a three-step empirical correction procedure inspired by the work of Leith [7], DelSole and Hou [8], and more recently Danforth et. al. [9, 13]. We first test its effectiveness in synchronizing Lorenz systems [10] with varied parameter-values in a perfect model scenario. We then apply the correction to an alternative model derived by Ehrhard and Müller [14] tuned to approximate the evolution of a toroidal thermosyphon, an experimental analogue to the original Lorenz system. The “true” climate is represented by a long-time, high-dimensional computational fluid dynamics (CFD) simulation of the thermosyphon. An *analysis*, which is an approximation of the true system state in terms of model variables, is then created by three-dimensional variational (3DVar) data assimilation and used for training and verification of the empirical correction. This process mimics the application of empirical correction in an operational NWP setting. We also verify the corrected model by direct comparison with observations of the truth.

The results in each experiment suggest that the correction procedure is effective for reducing error. However, there is an associated cost of this short-term error reduction, which is evidenced by substantial qualitative differences between the dynamics of the corrected model and the true system, differences that were not present in the uncorrected model. Introduction of system-specific knowledge into the correction procedure is shown to mitigate some of that cost, while also improving error statistics further than the entirely general procedure.

The paper is structured as follows. In Sec. II we define the procedure for a general model M . The application of the technique in the perfect model scenario is addressed in Sec. III, and in Sec. IV we present the thermosyphon model correction. Finally, we discuss the results and conclude the paper in Sec. V.

II. EMPIRICAL CORRECTION

The correction procedure employed in this experiment consists of three steps: (A) training, (B) state-independent correction and (C) state-dependent correction. The state-independent correction can be thought of as aligning the time-average of the model state with that of the true state. Likewise, the state-dependent correction can be considered an alignment of the model variance with that of the truth. To determine the correction

terms, we compare short model forecasts to observations of the true system over a training period in a process called *direct insertion*.

A. Training

In general, comparing model forecasts to a true physical system requires estimates of the true system state in terms of the model state-variables. Consider a vector time-series $\mathbf{x}^T(t)$ of such estimates, which we will call the “reference truth”. The amount of time h , measured in model timesteps, between estimates is called the *analysis window*; we assume it to be constant.

The process of direct insertion begins with the generation of a time-series $\mathbf{x}^M(t)$ of duration- h model forecasts, where each forecast in the time-series is initialized from the previous state in the reference truth. The first vector in the series, for example, will be $\mathbf{x}^M(t_0 + h)$, which is the model state resulting from an h -timestep forecast started with initial condition $\mathbf{x}^T(t_0)$. Subtracting each of the model forecast states $\mathbf{x}^M(t)$ from the corresponding reference true state $\mathbf{x}^T(t)$ produces a third time-series $\Delta\mathbf{x}(t)$ which represents the forecast errors after h timesteps. These errors result from differences between the model rate of change for each variable and the true rate of change, and they are commonly referred to as analysis corrections (or increments). See Fig. 1 for a schematic of the procedure.

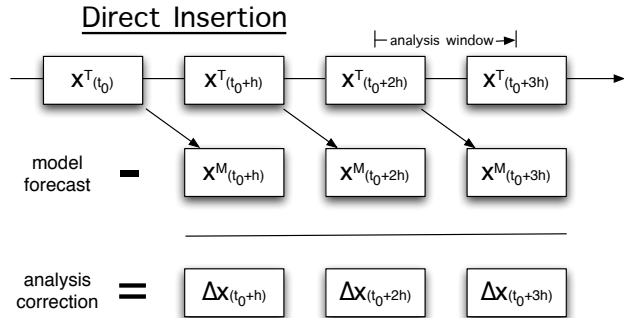


FIG. 1. The direct insertion procedure for comparing short model forecasts to the truth to obtain a time series of analysis corrections. \mathbf{x}^T represents a time series of the reference truth, and the analysis window represents the number of timesteps between estimations of the true system state. \mathbf{x}^M represents a time series of forecasts with duration equal to the analysis window, each of which is started from the previous true state. The time-average of the analysis corrections $\langle \Delta\mathbf{x} \rangle$ divided by the number of timesteps in the analysis window h approximates the average (state-independent) model bias \mathbf{b} .

Finally, we separate each of the time series into anoma-

lous and time-average components:

$$\begin{aligned}\mathbf{x}^{T'}(t) &= \mathbf{x}^T(t) - \langle \mathbf{x}^T \rangle \\ \mathbf{x}^{M'}(t) &= \mathbf{x}^M(t) - \langle \mathbf{x}^M \rangle \\ \Delta \mathbf{x}'(t) &= \Delta \mathbf{x}(t) - \langle \Delta \mathbf{x} \rangle\end{aligned}\quad (3)$$

where the expectation operator $\langle \cdot \rangle$ denotes averaging over the training period, and the primes denote *anomalies*, which are differences from the mean. The time-average components will be used for state-independent correction as described in Sec. II B and the anomaly time-series will be used for state-dependent correction as detailed in Sec. II C.

B. State-independent correction

We turn our attention first to a state-independent correction of the form

$$\dot{\mathbf{x}}^T \approx M^*(\mathbf{x}) \equiv M(\mathbf{x}) + \mathbf{b} \quad (4)$$

where the constant vector \mathbf{b} is the average model error (bias) to be determined.

Recall that our goal here is to empirically align the time-averages of the state-variables in the model to those of the true system. We call the time-averaged true system state the *climatology*, and we approximate it by $\langle \mathbf{x}^T \rangle$, the average of the reference true state over the entire training period. The average of the analysis corrections $\langle \Delta \mathbf{x} \rangle$ over the training period provides an estimate for the systematic, state-independent error generated by the model during the analysis window, as explained in Fig. 1. Dividing by the number of timesteps in the analysis window, then, we approximate the model bias by $\mathbf{b} = \langle \Delta \mathbf{x} \rangle / h$, and the bias-corrected model tendency is thus given by

$$M^*(\mathbf{x}) \equiv M(\mathbf{x}) + \frac{\langle \Delta \mathbf{x} \rangle}{h} \quad (5)$$

Note that at this point we are approximating the model tendency error δM by the model bias \mathbf{b} alone. We also wish to estimate any component of error that may depend on the system state, by approximating $\delta M \approx \mathbf{b} + \mathbf{L}\mathbf{x}'$, where \mathbf{L} is a matrix operator to be described in the next section.

C. State-dependent correction

To generate a linear state-dependent correction operator \mathbf{L} , we follow Leith [7] and DelSole and Hou [8], by first recomputing the forecast and correction time series in Fig. 1 using the state-independent corrected model, M^* , and then decomposing into mean and anomalous components. We seek an improved model of the form

$$\dot{\mathbf{x}}^T \approx M^+(\mathbf{x}) \equiv M^*(\mathbf{x}) + \mathbf{L}\mathbf{x}' \quad (6)$$

where M^+ includes both stages of correction. Letting $\mathbf{g} = \delta M^+ = \dot{\mathbf{x}}^T - [M^*(\mathbf{x}) + \mathbf{L}\mathbf{x}']$ be the tendency error of the improved model, we minimize the expected square tendency error $\langle \mathbf{g}^{tr} \mathbf{g} \rangle$, (where \mathbf{g}^{tr} is the transpose of \mathbf{g}), with respect to \mathbf{L} . The minimization results in the formula

$$\mathbf{L} = \langle \Delta \mathbf{x}' \mathbf{x}^{T' tr} \rangle \langle \mathbf{x}^{T'} \mathbf{x}^{T' tr} \rangle^{-1} \equiv \mathbf{C}_{\Delta \mathbf{x} \mathbf{x}^T} \mathbf{C}_{\mathbf{x}^T \mathbf{x}^T}^{-1} \quad (7)$$

as explained by Danforth et. al. [9]. $\mathbf{C}_{\Delta \mathbf{x} \mathbf{x}^T}$ is the average over the training sample of cross covariance matrices obtained by taking the outer product $[\Delta \mathbf{x}'(t)] \cdot [\mathbf{x}^{T'}(t)]^{tr}$ for each time t . $\mathbf{C}_{\mathbf{x}^T \mathbf{x}^T}$ is the average true-state covariance matrix, and \mathbf{L} is known as Leith's state-dependent correction operator.

When \mathbf{L} operates on the current anomalous state \mathbf{x}' , we can think of it as doing two things: (1) $\mathbf{C}_{\mathbf{x}^T \mathbf{x}^T}^{-1}$ operates on \mathbf{x}' , effectively relating the current state to the reference truth in the dependent sample, i.e. giving the best representation of the current state in terms of past states; and then (2) $\mathbf{C}_{\Delta \mathbf{x} \mathbf{x}^T}$ operates on the result, determining what correction should be made. This allows the model correction to adjust to different regions of state-space, and explains why the state-dependent correction can be thought of as attempting to align the model and true-state variances.

In the next section, we describe the application of this three-stage procedure to align Lorenz systems with different parameter values in a perfect model scenario, and in Sec. IV we describe its application to couple a low dimensional model to a high dimensional toy climate simulation. We also note here that the term *corrected* will imply the application of both state-independent and state-dependent correction, unless explicitly stated otherwise.

III. PERFECT MODEL SCENARIO

As a first step in the investigation of the correction technique, we consider its application to a model originally studied by Lorenz [10]. The system of equations (8) represents fluid flow between two plates, Rayleigh-Bénard convection, in which convection cells form for certain parameter ranges. However, with only slight modification (the details of which appear in Sec. IV A), they also describe the flow in a natural convection loop [14, 15]. Lorenz systems are covered exhaustively in publication [16, 17], and thus provide a familiar platform on which to perform preliminary tests of strategies for predicting the future state of chaotic systems.

A. Experimental design

In this perfect model scenario, the true system and the models share the structure of (8) and only differ in parameter values. Specifically, a true or *nature* run was

created by integrating the Lorenz system

$$\begin{aligned}\dot{x}_1 &= \sigma(x_2 - x_1) \\ \dot{x}_2 &= rx_1 - x_2 - x_1x_3 \\ \dot{x}_3 &= x_1x_2 - bx_3\end{aligned}\quad (8)$$

with the standard parameter set: $\sigma = 10$, $b = 8/3$, $r = 28$. Models with the same σ and b , but with r -values varying from 25 to 31 in increments of 0.5 (except for $r = 28$) were the subjects for correction. For each of these 12 models, the correction algorithm was performed using the 4 different analysis windows $h = 1, 2, 4$, and 8 timesteps, resulting in 48 distinct model-correction pairs in an exponential design. The training and testing of the corrected models is detailed in the following sections, and a picture showing one particularly positive outcome appears in Fig. 2.




FIG. 2. (Color online) Trajectories of the truth ($r = 28$), a model with r -perturbation -2 ($r = 26$), and the corrected model using an analysis window of $h = 8$ timesteps. All start from the same initial condition (circle), and represent 2 time-units (200 timesteps) of integration (concluding with squares). Note that the corrected model trajectory is well aligned with the true trajectory for much longer than the uncorrected model trajectory. Even after it deviates noticeably, the corrected model trajectory changes flow regimes (switches lobes) with the true trajectory. In contrast, the uncorrected model trajectory deviates from the true trajectory almost immediately, and remains in the initial lobe.

B. Training

A 100 time-unit nature run was generated by integrating system (8) from the initial condition $\mathbf{x}_0 = [1.508870, -1.531271, 25.46091]^{tr}$ (see [18]) for 10000 timesteps of $\kappa = 0.01$ time units each, using fourth order Runge-Kutta. For each of the 12 models and 4 analysis windows, a time-series of short forecasts was generated

by direct insertion (see Fig. 1). As an illustrative example, consider training with an analysis window of $h = 4$. The first forecast, $\mathbf{x}^M(t_0 + 4)$, is a 4-timestep model forecast started with the true initial condition $\mathbf{x}^T(t_0) = \mathbf{x}_0$. The second forecast, $\mathbf{x}^M(t_0 + 8)$, is a 4-timestep model forecast started with the true state $\mathbf{x}^T(t_0 + 4)$, and so on, resulting in $10000/4 = 2500$ total short forecasts. The state-independent correction $\langle \Delta \mathbf{x} \rangle / h$ was then computed as described in Fig. 1 and Sec. II B.

Next, the correction time-series $\Delta \mathbf{x}(t)$ was recomputed using a state-independent corrected version of the model, namely Eq. (5). Correction and true-state anomalies were calculated as in Eq. (3), and the cross covariance and covariance matrices were determined for each time t and averaged over the training sample to obtain $\mathbf{C}_{\Delta \mathbf{x} \mathbf{x}^T}$ and $\mathbf{C}_{\mathbf{x}^T \mathbf{x}^T}$ as outlined in Sec. II C. Finally, the static Leith operator \mathbf{L} was computed and the training procedure was complete.

Note that the training design imparts a statistical disadvantage upon the use of wider analysis windows. Specifically, doubling the analysis window halves the number of samples in the training period. The design was chosen, despite this prejudice, to more accurately reflect an operational implementation in which the training data is likely to be drawn from a fixed period of time. However, to further support the validity of comparisons between models corrected with different analysis windows, we note that letting the training period be $10000h$, ensuring that the number of samples is held constant at 10000, yields results that are qualitatively indistinguishable from those presented here.

C. Testing

A new nature run, 10000 time-units (one million timesteps) in length, was generated starting from the last true state in the training period. The purpose of beginning at the end of the training period was to obtain an *independent* truth with which to test the effectiveness of the corrected models. For each of the 48 corrected models, 1000 randomly selected states from this new nature run were used as the initial state, and both the uncorrected and corrected models were integrated for 20 time units. In addition, an *exact* model, with the same parameters as the truth, was integrated from a random perturbation of that same initial condition, on the order of 10^{-3} in each state variable. The purpose of the exact model is to represent an upper bound for the effectiveness of the empirical model error correction, i.e. demonstrating a case where the impact of the r -perturbation was corrected exactly, and forecast error comes only from the initial condition discrepancy.

Fig. 2 depicts the results of a single test with a particularly positive outcome. Trajectories of the truth (solid black), uncorrected model (blue dot), and corrected model (red dash), all start from the same initial state and progress for two time-units. The corrected




FIG. 3. (Color online) Plots of average RMSE over 1000 trials for the uncorrected (thick solid blue) and corrected models (all with $r = 26$) using analysis windows of 8, 4, 2 and 1, black dot, red dash-dot, green dash, solid magenta respectively. The exact model (thin solid black) is the same as the true system, but started from an initial condition perturbed randomly on the order of 10^{-3} in each state variable. Time units are on the x -axis, where one time unit is 100 timesteps in the numerical integration. The technique's performance clearly degrades with widening analysis window.

model trajectory stays close to the true trajectory for much longer than the uncorrected model trajectory, and even after deviating markedly, the corrected model trajectory still switches lobes with the true trajectory.

Two metrics were used to measure forecast accuracy: root mean square error (RMSE), and anomaly correlation (AC). The RMSE is given at time t by

$$\text{RMSE}(t) = \sqrt{\frac{1}{N} \sum_{n=1}^N [x_n^T(t) - x_n^M(t)]^2} \quad (9)$$

where \mathbf{x}^M is the model state and \mathbf{x}^T is the true state, and we are summing over the $N = 3$ state variables in the Lorenz system. Fig. 3 plots the average RMSE over 1000 trials performed for correction of the $r = 26$ (r -perturbation of -2) model using the 4 different analysis windows.

Anomaly correlation is a metric frequently used in weather and climate modeling to determine the length of time for which a model forecast is useful. The AC is given by

$$\text{AC} = \frac{\mathbf{x}^{M'} \cdot \mathbf{x}^{T'}}{\|\mathbf{x}^{M'}\|_2 \|\mathbf{x}^{T'}\|_2} \quad (10)$$

where $\mathbf{x}^{M'}$ and $\mathbf{x}^{T'}$ are the anomalous model state and anomalous true state, respectively, at a particular time. AC is essentially the dot product of the anomalous model state with the anomalous true state, normalized such that

AC = 1 for a perfect model. A forecast is typically considered useful for as long as its AC remains above 0.6 [6]. As with RMSE, the AC scores for each model, corrected and uncorrected, were averaged over 1000 trials to provide a good representation of model performance. See Fig. 4 for AC plots demonstrating the effects of changing analysis window length and parameter perturbation in the original model on the duration of useful forecasts.

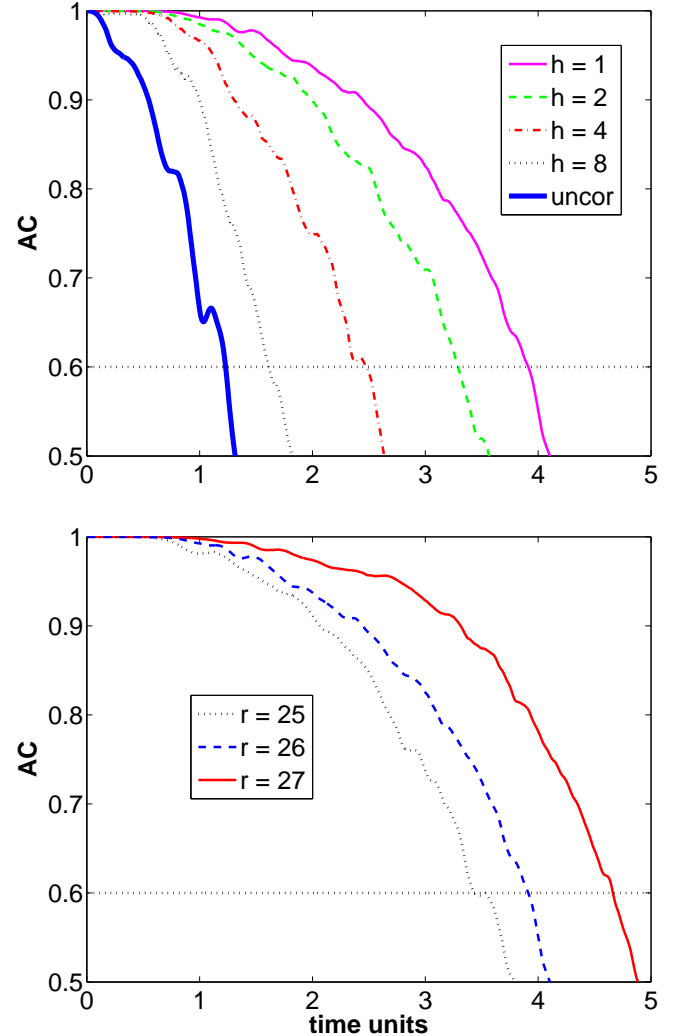


FIG. 4. (Color online) (top) Plots of average AC over 1000 trials for the uncorrected (thick solid blue) and corrected models (all with $r = 26$) using analysis windows of $h = 8, 4, 2$ and 1, black dot, red dash-dot, green dash, solid magenta respectively. In agreement with Fig. 3, performance drops as the analysis window widens. Note, however, that even with $h = 4$ timesteps the corrected model provides a useful forecast for twice as long as the uncorrected model. (bottom) Plots of average AC for models corrected with an analysis window of 1 timestep, and $r = 25, 26$ and 27, black dot, blue dash, and solid red respectively. The greater the magnitude of r -perturbation, the more difficult it is to correct successfully. Similar results are observed for r -values greater than 28.

D. Results

Fig. 3 demonstrates that the tested empirical correction technique succeeds in reducing error in the perfect model scenario. However, the importance of frequent observations of the truth during the training period is highlighted by the approach of the average corrected model RMSE towards the uncorrected model RMSE with widening of the analysis window. It is also noteworthy that the corrected model remains almost as good as the exact model for a full time unit when observations are made every timestep. In this case the RMSE remains below 2 more than 5 times as long the uncorrected model.

Fig. 4 demonstrates that along with reduced error, empirical correction has the potential to provide forecasts that are useful for much longer. Training with an analysis window of 1 timestep, the corrected model forecasts are useful for nearly 4 times longer than the uncorrected model forecasts. In light of the sensitivity of AC to analysis window length, the bottom panel of Fig. 4 suggests that the accuracy of parameter values matters less for the effectiveness of the corrected model, as measured by error statistics, than does the frequency of observations in training. However, it should be noted that error statistics are not the whole story. The ability of the corrected models to reproduce the qualitative dynamical *behavior* of the true system accurately, like switching of lobes in the attractor, is only indirectly indicated by reduced forecast error. We address this issue in Sec. IV D.

A summary representation of the data in Fig. 4 is shown in Fig. 5, where the duration of forecast usefulness is plotted vs. analysis window in the top panel, and vs. r -perturbation in the bottom panel. However, these plots still focus only on two cross sections of the 48 total model-correction pairs that were tested in the perfect model scenario, specifically the 4 corrected $r = 26$ models in the top and the 12 corrections across all r -perturbations using an analysis window of 1 timestep in the bottom. To get a better sense of the relative importance, Fig. 6 shows a surface plot of duration of useful forecasts for all combinations of analysis window and r -perturbation tested. Surprisingly, with an analysis window of 2 timesteps, the corrected models with the largest r -perturbation (parameter error greater than 10%) produce forecasts that are useful longer than those made by the uncorrected model with the smallest r -perturbation (parameter error less than 2%). For systems with reasonably small model errors, this is an indication that empirical correction may improve forecasts more readily than parameter tuning.

IV. TOY CLIMATE MODEL

We next investigate the effectiveness of the correction procedure in a more realistic situation, where the forecast model is structurally different (in dynamics, dimension, parameterization, etc...) from the true system. Consider




FIG. 5. (Color online) Plots of the duration of useful forecasts (first time AC=0.6) vs. analysis window (top) and r -perturbation (bottom). (top) The drop in duration of useful skill with analysis window suggests an asymptotic decrease in the value of correction towards no improvement from the uncorrected $r = 26$ model. (bottom) Using an analysis window of 1 timestep, we plot the duration as a function of r -perturbation. The black diamond is the average duration of a useful forecast for the exact model (true system with perturbed initial state), representing the limit of predictability imposed by initial condition uncertainty. Note: the corrected models with greatest r -perturbation outperform the uncorrected models with least r -perturbation.

a fluid-filled, vertically-oriented natural convection loop, or thermosyphon, with circular geometry, a schematic of which appears in Fig. 7. The constant temperature imposed on the wall of the lower half of the loop, T_h , is greater than the constant temperature imposed on the wall of the upper half, T_c , resulting in a temperature inversion. For large enough temperature differences $\Delta T_w = T_h - T_c$, convection dominates, and the flow undergoes chaotic reversals of direction referred to as *flow regime changes*, while remaining laminar [14, 15]. These dynamics produce forecasting difficulties similar to those




FIG. 6. (Color online) Surface plot showing duration of useful forecasts for all combinations of analysis window and r -perturbation tested, top surface corrected, bottom surface uncorrected. On the left face we can see the cross section represented by Fig. 5(b). Using an analysis window of 2 timesteps, corrected models with largest r -perturbation provide forecasts that are useful longer than those of uncorrected models with smallest r -perturbation.

encountered in weather and climate prediction, and thus the thermosyphon provides a useful platform on which to test potential improvements to forecasting methods.

A. Experimental design

The true system was represented by numerical simulation using the 2-D laminar Navier-Stokes equations along with the energy equation, and a finite-volume-based flow modeling software package (FLUENT 6.3) was used to perform the numerical integration, see [15] for details. Almost 90 days of fluid behavior was generated, in which $O(10^4)$ flow reversals occurred. We also note that for the Rayleigh number used in this experiment, $Ra = 1.5 \times 10^5$, the thermosyphon has two unstable convective equilibrium solutions corresponding to steady clockwise and steady counter-clockwise flow [15, 19].

The low-dimensional model used to make forecasts of the CFD simulation is the Ehrhard-Müller (EM) system:

$$\begin{aligned} \dot{x}_1 &= \alpha(x_2 - x_1) \\ \dot{x}_2 &= \beta x_1 - x_2(1 + KH(|x_1|)) - x_1 x_3 \\ \dot{x}_3 &= x_1 x_2 - x_3(1 + KH(|x_1|)) \end{aligned} \quad (11)$$

where x_1 is proportional to the mean fluid velocity, x_2 is proportional to the horizontal temperature difference in the loop, and x_3 is proportional to the deviation of the vertical temperature difference from the value it takes during conduction. This system was derived from physical principles to model a natural convection loop [14, 15],




FIG. 7. Schematic of the natural convection loop, or thermosyphon. The top half of the fluid-filled loop is cooled to the temperature T_c , while the bottom half is heated to $T_h > T_c$, creating a temperature inversion. The three state variables in the low-dimensional model are (1) $x_1 \propto \bar{u}$, the mass flow rate (mean fluid velocity); (2) $x_2 \propto \Delta T_{3-9}$, the horizontal temperature difference; and (3) $x_3 \propto \Delta T_{6-12} - \Delta T_{6-12}^{\text{cond.}}$, the difference between the vertical temperature profile and the value it takes during conduction.

and for this study the parameters $\alpha = 7, \beta = 33, K = 0.07$ were tuned empirically to best match the flow reversal behavior of the CFD simulated thermosyphon.

The primary difference between EM and Lorenz systems is H , a function that determines the velocity dependence of the heat transfer between the fluid and the wall. This characteristic of the flow is ignored by the Lorenz equations (i.e. $K = 0$). H varies as the third root of the magnitude of the mean fluid velocity for $|x_1| > 1$, and as a fourth degree polynomial in $|x_1|$ for $|x_1| \leq 1$ to remain differentiable; the reader may see [15] for more detail. We note that when $K = 0$ in the EM equations (11), they are identical to the Lorenz system (8) with $b = 1$. Physically, the unitary geometric factor (i.e. $b = 1$) in EM results from the forced single circular convection cell in a thermosyphon, as opposed to the unconstrained flow producing multiple cells between two plates.

Using a background forecast created with the EM model, and observations of the CFD mean fluid velocity \bar{u} with Gaussian noise added to simulate error, 3DVar data assimilation was performed to generate an analysis, or best guess of the true state of the system in terms of the variables of the forecast model [15]. One segment of the analysis is used as the reference truth $\mathbf{x}^T(t)$ for training, and the remainder is used for testing.

B. Training

Approximately 3 days of 3DVar analysis, corresponding to 432 time-units in the EM forecast model, was used as the training period reference truth. As in the perfect model scenario, fourth-order Runge-Kutta was used with a timestep of $\kappa = 0.01$ time-units. An analysis window of $h = 5$ timesteps, corresponding to about 30 seconds of simulated flow, was used to match the frequency of observation in the data assimilation scheme. Thus, 8640 short forecasts were used to compute the model bias and Leith operator for the model, as outlined in Sec. II.

C. Testing

Three forecast models were compared by verification against both the 3DVar analysis and direct observation of the mass flow-rate \bar{u} in the CFD simulated thermosyphon: (1) the uncorrected EM model with parameters tuned to best represent observations of the CFD simulated mass flow rate; (2) the tuned model with correction applied; and (3) an EM model whose parameters differ from the tuned model by 10%, *with* correction applied. The purpose of the third test-model is to gauge the relative capabilities of empirical correction and parameter tuning for error reduction and prolonging the usefulness of forecasts.

A set of 1000 trial forecasts were performed, starting from randomly chosen points in the analysis *after* the training period. For each forecast, RMSE and AC time-series were computed with respect to the analysis and averaged over all trials. See the top panel of Fig. 8 for the resulting average AC plot (RMSE not shown). We also verify the model forecasts against observed scalar mass flow-rate \bar{u} , for which time-series of relative error were averaged over the 1000 trials, pictured in the bottom panel of Fig. 8.

Two details are important in the computation of the relative error. First, to compare model output to observations, it is necessary to convert the model state-variables to “observation-space” variables. In other words, an observation operator determined by data assimilation was used to convert the model state-vector $(x_1, x_2, x_3)^{tr}$ to an observation-space value \tilde{x}_1 , which is the predicted mass flow-rate of the system. Second, the error is taken relative to the saturation point, which we define as the average absolute difference between the mass flow-rates of the system at randomly chosen points in time. Thus, an average relative error near 1 means that the forecast model is no better than a random guess.

The results presented in Fig. 8 indicate that corrected models, tuned or not, produce smaller short-term forecast error on average than the uncorrected, optimally tuned model. Corrected model forecasts are thus typically useful for longer. Though this is an important benefit, average short-term error statistics may conceal considerable qualitative differences between model dynamics




FIG. 8. (Color online) Plots showing verification of the corrected model with respect to (top) 3DVar analysis and (bottom) direct observation of the CFD simulated mass flow-rate \bar{u} . (top) We see that empirical correction produces forecasts that are useful approximately 25% longer than forecasts of the uncorrected model. Also, correction of an untuned model (10% difference from optimal value in every parameter) produces forecasts that are useful far longer than the uncorrected but optimally tuned model. (bottom) The corrected model shows reduced error when verified against direct observations of the mean fluid velocity in the simulation as well. Again we see that correction of a model with sub-optimal parameterization produces better results than tuning the parameters optimally but neglecting to apply empirical correction.

and those underlying the true system. Stability of equilibrium solutions, and changes of flow regime characterized by aperiodic switching between otherwise confined regions of state-space are examples of qualitative characteristics for which it may be crucial that the model dynamics match those of the truth. In the next section we address the effect of empirical correction on the dynamical matching capability of the EM forecast model.

D. Matching qualitative dynamics

To measure the accuracy of models with regard to matching the flow reversal behavior of the true system, forecasts were generated with both the corrected and uncorrected EM models for 5000 initial states throughout the attractor, from the testing portion of the 90-day 3DVar analysis, and the time of the first predicted flow reversal was recorded for each one. We investigate the difference between the predicted times and the actual times (from the analysis) of first flow reversal, taken $t_{\text{model}} - t_{\text{actual}}$, so that positive values indicate late predictions while negative values indicate early predictions. See Fig. 9 for plots of the results.

Three unforeseen costs of empirical correction for this system, in terms of lost dynamical matching capability, can be summarized as follows:

- 1) Stabilization of convective equilibrium solutions
- 2) Elimination of flow reversal behavior for states in a neighborhood of either convective equilibrium
- 3) Spurious dynamical asymmetry between lobes

We address these costs in the sections that follow, examining the first two related phenomena in Sec. IVD 1, and then explaining the spurious asymmetry in Sec. IVD 2.

1. Stabilization of equilibrium solutions

Jacobian analysis of the EM equations (11) provides analytical confirmation of the instability of the convective equilibria in the uncorrected model. Specifically, the Jacobian evaluated at each equilibrium has one negative real eigenvalue, whose eigenvectors are in the local direction of (tangent to) the stable manifold of the equilibrium, and a conjugate pair of complex eigenvalues with positive real part, whose 2-D eigenspace is locally tangent to the unstable manifold of the equilibrium. In fact, for both convective equilibria the positive real parts of these unstable eigenvalues are quite small, on the order of 10^{-2} , indicating weakly repelling instability. In the following we explain analytically the mechanism by which empirical correction overcomes this weak repulsion, producing a forecast model with *attracting*, and thus stable, convective equilibria, see Fig. 10.

Empirical correction of the EM model effectively alters the right-hand side of the differential equations (11) by first adding a constant related to the bias term \mathbf{b} , and then adding a term that depends linearly on the model state, i.e., something related to $\mathbf{L}\mathbf{x}'$. Letting \mathbf{f}^M be the vector-valued EM differential equation, and \mathbf{f}^{M+} be the corrected equation, we write

$$\mathbf{f}^{M+} = \mathbf{f}^M + \mathbf{b}_0 + \mathbf{L}_0\mathbf{x}' \quad (12)$$

where

$$\mathbf{b}_0 = \lim_{\kappa \rightarrow 0} \mathbf{b} \quad \text{and} \quad \mathbf{L}_0 = \lim_{\kappa \rightarrow 0} \mathbf{L} \quad (13)$$

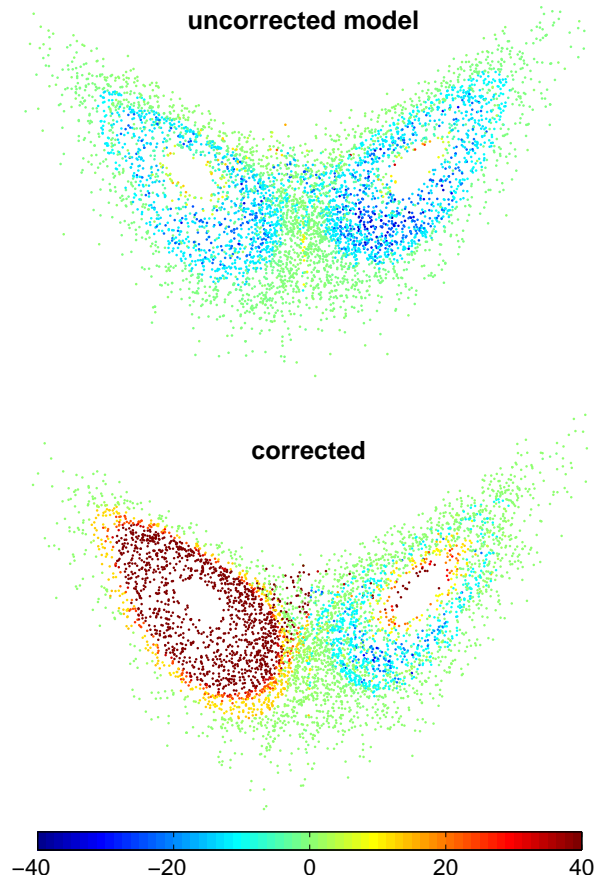


FIG. 9. (Color online) Difference (in minutes) between predicted and actual time of first flow reversal, plotted by initial state, for the uncorrected (top) and corrected (bottom) models. The difference was taken $t_{\text{model}} - t_{\text{actual}}$, so that positive values (towards red) indicate late predictions while negative values (towards blue) indicate early predictions. First, note that for initial states near the edge of the attractor, in either lobe, each model performs well in predicting the coming flow reversal. This makes sense because for these initial states the change of flow regimes is imminent, and thus easily predicted. Next, observe that when the uncorrected model errs, it almost always makes an early prediction (blue dots). Model trajectories close to the unstable convecting equilibrium solution at the center of each lobe oscillate with larger amplitude than in the true system. One might expect that empirical correction would counteract the effect, and in fact that is apparent (bottom). However, rather than improving flow reversal prediction, the results here indicate extreme overcorrection. The dark red dots are not 40-minute late predictions; the limit of the color axis was chosen to better illustrate the spread in the remaining data. In actuality, the corrected model predicts that a flow reversal *will never happen* for those initial states. Finally, observe that the empirical correction has manufactured a lobe asymmetry, characterized by the much smaller region of dark red in the right lobe.

can be thought of as the computed bias term and Leith operator, respectively, for an infinitesimal timestep κ . Because of the nonlinearity of the system, we cannot determine the exact relationship between the infinitesimal correctors \mathbf{b}_0 and \mathbf{L}_0 , and the bias term \mathbf{b} and Leith operator \mathbf{L} , respectively, that we compute using a timestep of $\kappa = 0.01$ and analysis window of $h = 5$ timesteps. However, since we discretize \mathbf{f}^{M^+} in numerical integration, we *can* determine the \mathbf{b}_0 and \mathbf{L}_0 that we actually apply. We effectively approximate the correction terms

$$\mathbf{b}_0 \approx \frac{\mathbf{b}}{\kappa} = 100\mathbf{b} \quad \text{and} \quad \mathbf{L}_0 \approx \frac{\mathbf{L}}{\kappa} = 100\mathbf{L} \quad (14)$$

within the fourth order Runge-Kutta scheme.

Now, armed with an analytical representation of the differential equations \mathbf{f}^{M^+} , we note the following relationship between the corrected model Jacobian $D\mathbf{f}^{M^+}$, and that of the uncorrected EM model $D\mathbf{f}^M$:

$$D\mathbf{f}^{M^+} = D\mathbf{f}^M + \mathbf{L}_0 \approx D\mathbf{f}^M + 100\mathbf{L} \quad (15)$$

since the constant bias term \mathbf{b}_0 disappears and \mathbf{L}_0 operates on a translation of the model state. We evaluate the Jacobian of the corrected model at each of the convective equilibria, and determine its eigenvalues. Indeed, the real part of the complex conjugate eigenvalues, for each convective equilibria, is *negative* for the corrected model \mathbf{f}^{M^+} , and thus the convective equilibria have become stable.

The equilibria attract all states inside neighborhoods around them, which thereby separate state-space near the attractor into regions whose trajectories will either change flow regime at least once, or not at all. See Fig. 10 for an example of this effect. Furthermore, any trajectories that land in one of these neighborhoods after only one or two flow reversals, which might occur within the expected 17-minute duration of useful forecasts for the corrected model reported in Fig. 8, will then approach steady convection. This behavior is in qualitative *opposition* to that of the true system, *and* that of the original uncorrected EM model, for which steady convection in a single direction is an unstable equilibrium.

2. Broken symmetry

The size discrepancy between left and right-lobe regions attracted to the convective equilibria of the corrected model revealed in Fig. 9 demonstrates that empirical correction breaks the symmetry of the EM system. As in the conventional Lorenz system (8), the EM system (11) is symmetric under the mapping $(x_1, x_2, x_3) \mapsto (-x_1, -x_2, x_3)$. Again letting \mathbf{f}^M be the vector-valued EM differential equation, this symmetry implies that \mathbf{f}^M commutes with a certain matrix A , i.e.

$$A\mathbf{f}^M(\mathbf{x}) = \mathbf{f}^M(A\mathbf{x}), \quad A = \begin{bmatrix} -1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (16)$$

In fact, empirical correction breaks this symmetry in two ways.

First, recall that after bias correction alone we have changed the EM system by adding a constant vector to the right-hand side of the differential equation. Letting \mathbf{f}^{M^*} represent the bias-corrected differential equation, \mathbf{f}^{M^*} *does not* commute with A unless there is zero bias in x_1 and x_2 , i.e.

$$A\mathbf{f}^{M^*}(\mathbf{x}) \neq \mathbf{f}^{M^*}(A\mathbf{x}) \quad \forall \mathbf{b}_0 \neq \begin{bmatrix} 0 \\ 0 \\ b_3 \end{bmatrix} \quad (17)$$

where b_3 can be any constant, and we recall that $\mathbf{b}_0 = \lim_{\kappa \rightarrow 0} \mathbf{b}$ can be thought of as the computed bias term for an infinitesimal timestep κ . Note that even if no bias in x_1 or x_2 existed, the probability of statistically computing a bias term \mathbf{b} that would preserve the symmetry of the EM system is zero.

In the unlikely case that a bias term is computed that preserves symmetry, or such a bias term is forced, state-dependent correction will break it. Assuming that \mathbf{f}^{M^*} does commute with A , and letting \mathbf{f}^{M^+} be the fully corrected differential equation, then

$$A\mathbf{f}^{M^+}(\mathbf{x}) = \mathbf{f}^{M^+}(A\mathbf{x}) \iff A\mathbf{L}_0 = \mathbf{L}_0A \quad (18)$$

In other words, \mathbf{f}^{M^+} commutes with A if and only if $\mathbf{L}_0 = \lim_{\kappa \rightarrow 0} \mathbf{L}$ commutes with A . This forces the computed \mathbf{L} to be of the form

$$\mathbf{L} = \begin{bmatrix} \ell_{11} & \ell_{12} & 0 \\ \ell_{21} & \ell_{22} & 0 \\ 0 & 0 & \ell_{33} \end{bmatrix} \quad (19)$$

where the ℓ_{ij} can be any constants. Of course the probability of computing such an \mathbf{L} statistically is also zero. Therefore, both bias and state-dependent correction break the symmetry of the EM model.

3. Summary

Undesirable effects of empirical correction include the breaking of system symmetry, and altered stability of equilibrium solutions which results in large regions of state-space for which flow-reversals never occur in corrected model forecasts. We emphasize that the qualitative behavior of the corrected model is substantively different from that of the uncorrected model, which matches the behavior of the CFD simulated truth. This is true despite the fact that the corrected model shows improved average error statistics. However, it is possible to adjust the correction procedure to mitigate this effect by directly incorporating dynamical knowledge of the true system, which is the subject of the next section. Note that in doing so, we sacrifice the general applicability of the technique.




FIG. 10. (Color online) The first 20 minutes of a forecast for an initial state attracted to the left-lobe convective equilibrium stabilized by the empirical correction (i.e., one of the dark red dots towards the center of the left lobe in the bottom of Fig. 9). The uncorrected EM model trajectory (blue) behaves as it should, winding away from the equilibrium at the center of the lobe, whereas the trajectory of the corrected model (red) collapses toward the equilibrium solution, indicating the false stability produced by the correction. The source of the false stability is the overcorrective nudging visible as jumps (every 30 seconds, the length of the analysis window) on the red curve. The correction is attempting to lengthen flow regimes, a consequence of the blue regions near the equilibria in the top of Fig. 9.

E. Incorporating dynamical knowledge

To encode dynamical knowledge of the system in the empirical correction procedure, the state-space is partitioned into regions based on the qualitative behavior of the system, and then a separate bias term \mathbf{b} and state-dependent correction operator \mathbf{L} is computed for each region. For example, in the context of weather forecasting, the state-space of the atmosphere could be divided by stage in the El Niño oscillation, or by day and night, or local season for regional models. Fig. 9 suggests two ways to partition the state-space in the present context: (1) by flow regime direction (lobe); or (2) by distance from the nearest convective equilibrium solution. In each case, the state-space is decomposed into two regions, left/right lobe, or near/far from equilibrium, respectively. In addition to testing the correction procedure using each of these strategies individually, a procedure applying them simultaneously, which results in a partition of the state-space into four regions, is also tested.

1. Lobe dependence

To generate lobe-dependent bias correction terms and state-dependent correction operators, two regions $L_{1,2}$, corresponding to flow regimes of opposite direction, are defined by

$$L_1 = \{\mathbf{x} : x_1 < 0\}, \quad L_2 = \{\mathbf{x} : x_1 \geq 0\} \quad (20)$$

noting that their union is the entire state-space. Physically, L_1 (L_2) represents all states undergoing clockwise (counter-clockwise) convection. Next, the direct insertion procedure illustrated in Fig. 1 is modified to produce two subsequences of the analysis correction time series, $\Delta\mathbf{x}_{L_1}$ and $\Delta\mathbf{x}_{L_2}$, that correspond to two subsequences of the analysis time series, $\mathbf{x}_{L_1}^T$ and $\mathbf{x}_{L_2}^T$, respectively, where

$$\mathbf{x}_{L_1}^T = \{\mathbf{x}^T(t) : \mathbf{x}^T(t) \in L_1\} \quad (21)$$

and $\mathbf{x}_{L_2}^T$ is defined similarly. These subsequences are separated into mean and anomalous components, as in Eqs. (3), using means over each individual subsequence. Finally, the separate correction terms and operators are computed by:

$$\mathbf{b}_{L_k} = \frac{\langle \Delta\mathbf{x}_{L_k} \rangle}{h}, \quad \mathbf{L}_{L_k} = \mathbf{C}_{\Delta\mathbf{x}_{L_k} \mathbf{x}_{L_k}^T} \mathbf{C}_{\mathbf{x}_{L_k}^T \mathbf{x}_{L_k}^T}^{-1} \quad (22)$$

for $k = 1, 2$. To apply the lobe-dependent correction, at every timestep of numerical integration the current state is determined to be in either L_1 or L_2 , and the appropriate bias correction term and state-dependent operator are applied to advance the model.

2. Dependence on distance from equilibrium

A procedure analogous to that for lobe-dependent correction is applied here. Two regions $E_{1,2}$, corresponding to near and far from equilibrium, respectively, are defined by

$$\begin{aligned} E_1 &= \{\mathbf{x} : \min(\|\mathbf{x} - \mathbf{eq}_1\|_2, \|\mathbf{x} - \mathbf{eq}_2\|_2) < c_v\}, \\ E_2 &= \{\mathbf{x} : \mathbf{x} \notin E_1\} \end{aligned} \quad (23)$$

where \mathbf{eq}_1 and \mathbf{eq}_2 are the convective equilibria (estimated from the parameters of uncorrected model, [15]), and the critical value c_v is a parameter for the procedure. We are effectively approximating the neighborhoods attracted to the convective equilibria by spheres of radius c_v . For all results shown in the paper, the critical value $c_v = 8.5$ was used, though error statistics and dynamical matching capability were virtually unaffected by changing this parameter by 25% in either direction. This range of critical values was tested as estimates to the average radius of the dark red region in the left lobe of Fig. 9 (bottom). Continuing with the correction scheme, direct

insertion is modified as in the lobe-dependent correction, and the region-specific bias terms and state-dependent operators are calculated as in Eqs. (22), substituting E_k for L_k . Application of the correction to a forecast model also proceeds in the same fashion.

3. Simultaneous lobe and equilibrium dependence

Defining the lobe regions $L_{1,2}$ as in the lobe-dependent section, and the equilibrium regions $E_{1,2}$ as in the previous section, we define the four regions for simultaneous lobe and equilibrium-dependent correction by

$$\begin{aligned} R_1 &= L_1 \cap E_1, & R_2 &= L_1 \cap E_2 \\ R_3 &= L_2 \cap E_1, & R_4 &= L_2 \cap E_2 \end{aligned} \quad (24)$$

so that we modify direct insertion to produce four subsequences of analysis increments, each paired with the appropriate subsequence of the analysis time series. Note that the critical value c_v used in defining the regions $E_{1,2}$ does not depend on the lobe in this scheme. Allowing a different c_v for each lobe is a possible modification that was not tested. We compute the bias terms and Leith operators as in Eqs. (22), substituting R_k for L_k , where now $k = 1, 2, 3, 4$. Again, application of the correction proceeds as in the individual dynamically informed schemes, where the current state is determined to be in one of the four defined regions, and the appropriate bias term and Leith operator are used to advance the model.

F. Results

First, we discuss the results of applying the dynamically uninformed, and entirely general correction procedure as described in Sec. II, to couple the three-dimensional EM forecast model to the high-dimensional CFD simulated thermosyphon. Then we report the results of modifying the correction procedure to directly incorporate dynamical knowledge of the system, as detailed in Sec. IV E.

1. General procedure

The ability of the correction procedure to overcome parameter inaccuracies, as suggested by the results in the perfect model scenario, inspired the comparison of three forecast models.

- (1) tuned, uncorrected
- (2) untuned, corrected
- (3) tuned and corrected

Fig. 8 shows that empirical correction produces forecasts that remain useful longer, and demonstrate reduced error in this mock-operational setting, as well as in the

perfect model scenario. Verification against both analysis and observations allows more confidence in the success of the procedure in aligning forecasts made by the low-dimensional model with the CFD simulated thermosyphon.

Additionally, the results constitute evidence that the correction procedure is more effective than fine-tuning of parameters for improving error statistics in this more realistic setting, as well as in the perfect model scenario. This is not to say that parameters should not be tuned, but rather that empirical correction is a cheaper and more effective avenue for reducing forecast uncertainty in the present context.

Empirical correction of the EM forecast model comes at the cost of decoupling qualitative dynamics from those of the true system, however. As shown in Fig. 9, a large region is created in which initial states have trajectories under the forecast model that behave completely different from how they would in the true system. An example of such a trajectory appears in Fig. 10. While the general empirical correction procedure reduces forecast error for many initial states, it produces a forecast model that is entirely useless for others. By introducing dynamical knowledge into the correction procedure, the results of which appear in the next section, we greatly reduce the number of such initial states, and also further reduce average forecast error.

2. Dynamically informed correction

In an effort to produce forecast models that more accurately reflect the qualitative dynamics of the true system, two types of system-specific dynamical knowledge were directly incorporated into the empirical correction procedure: (1) flow regime direction of the current system-state, and (2) distance between the state and the nearest convective equilibrium solution. These dynamical cues were incorporated individually, and also simultaneously, resulting in three dynamically informed, empirically corrected forecast models, as outlined in Sec. IV E.

Fig. 11 shows average AC and mass flow-rate relative error over 5000 trials for the three dynamically informed and corrected models, as compared to the original biased model and the dynamically uninformed, corrected model. As described in the caption to the figure, encoding the current flow regime into the correction procedure results in a forecast model that is useful for almost twice as long as the original biased model, and doubles the improvement that was gained by dynamically uninformed empirical correction. Encoding the distance to the nearest equilibrium solution, on the other hand, does not greatly prolong usefulness, and in fact reduces it slightly when applied simultaneously with lobe-dependent correction.

In Fig. 12, we see how the three forecast models resulting from dynamically informed empirical correction compare with the dynamically uninformed, corrected model, with regard to matching the qualitative behavior of the




FIG. 11. (Color online) Average AC (top) and mass flow-rate relative error (bottom) over 5000 trials for models that were: uncorrected (thick solid blue), corrected (green dash-dot), equilibrium-dependent (ED) corrected (cyan dash), lobe and equilibrium-dependent (LD-ED) corrected (black dot), and lobe-dependent (LD) corrected (solid red). Note that the models are listed in order of increasing performance, as measured by these error statistics. Lobe-dependent correction, i.e., correcting with bias terms and state-dependent operators that depend on the current flow regime direction, produces forecasts that are useful for almost twice as long as those made by the uncorrected model, on average. When compared to the dynamically uninformed correction scheme (green dash-dot), lobe-dependent correction more than doubles the improvement over the uncorrected model. We note that equilibrium-dependent correction does not produce much improvement in forecast error over the dynamically uninformed procedure, and hinders the performance when applied simultaneously with lobe-dependent correction.

true system. We see a trend of improvement, characterized by smaller regions of states whose dynamics are different in the models than in the CFD simulated thermosyphon (dark red regions), as we apply first lobe-dependent, then equilibrium-dependent, and finally simultaneously lobe and equilibrium-dependent correction.

Results shown in Figures 11 and 12 suggest that encoding flow regime direction into the correction proce-

dures primarily enhances forecast statistics, while encod-

TABLE I. Median absolute differences between predicted and actual times of first flow reversal (row 1), along with percentage of trials for which the first flow reversal was predicted within 1 and 2 minutes, (rows 2 and 3, respectively), for the uncorrected model (M), dynamically uninformed corrected model (CM), lobe-dependent corrected model (LD), equilibrium-dependent corrected model (ED), and simultaneously lobe and equilibrium-dependent corrected model (LD-ED).

	M	CM	LD	ED	LD-ED
median	4	10	1.5	10.5	1.5
% within 1min.	38.26	38.54	49.12	38.68	48.64
% within 2min.	47.04	43.40	54.38	43.58	57.08

ing distance from equilibrium primarily enhances dynamical matching. Simultaneous inclusion of the two types of dynamical cues results in the best dynamical matching, with only a slight cost in average forecast accuracy. For further evidence of this summary conclusion, consider Table I.

In the first row, the median absolute differences between predicted and actual times of first flow reversal over the 5000 trials are listed for the uncorrected model (M), dynamically uninformed corrected model (CM), lobe-dependent corrected model (LD), equilibrium-dependent corrected model (ED) and simultaneously lobe and equilibrium-dependent model (LD-ED). The means are highly skewed for the CM, LD and ED models, due primarily to the large basins of attraction for the left-lobe equilibrium solution, see Fig. 12 top left, right and bottom left, respectively, and thus we show only the medians. The medians show that both the LD and the LD-ED corrected models predict the first flow reversal more accurately than the uncorrected model. The bottom two rows show the percentage of the 5000 trials for which the first flow reversal was predicted within 1 and 2 minutes, respectively, for each of the models. Again the LD and LD-ED corrected models show the best performance. Note that the LD-ED model boosts both the 1 and 2-minute success rates by approximately 10% over the uncorrected model.

We note, however, that even the LD-ED model exhibits a spuriously stable convective equilibria in each lobe, Fig. 12. Incorporating dynamical knowledge of the true system in the correction procedure, at least through the partitioning of state-space as we have done here, is not enough to avoid the stabilizing effect of empirical correction on the equilibria of the EM model explained in Sec. IV D 1. It is plausible that models of other systems with weakly repelling (attracting) equilibria might be subject to similar stabilizing (destabilizing) effects under empirical correction. Such effects might be mitigated by the strategy employed in the present work, but probably not avoided.

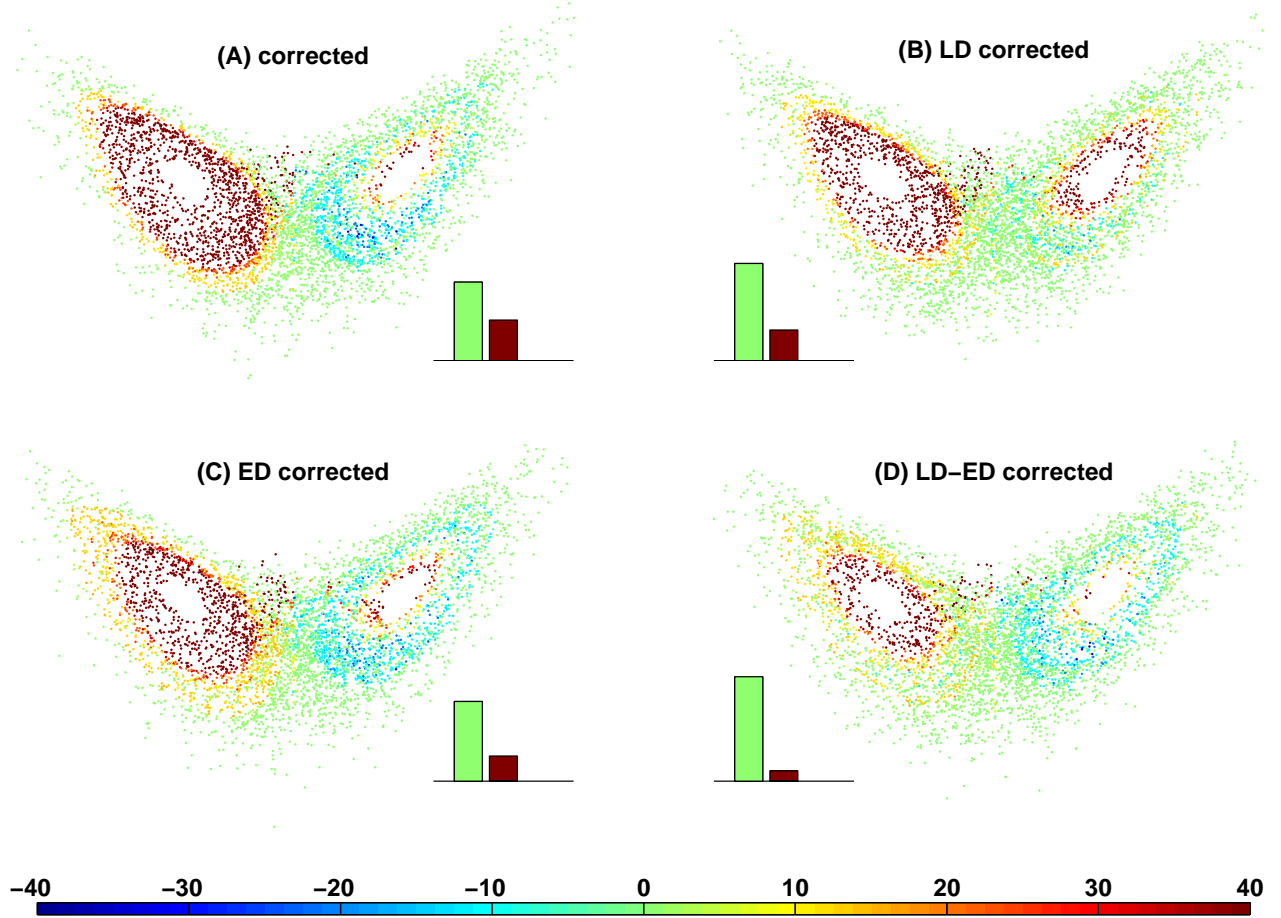


FIG. 12. (Color online) Difference (in minutes) between predicted and actual time of first flow reversal, plotted by initial state, for the (A) corrected, (B) lobe-dependent corrected, (C) equilibrium-dependent, and (D) lobe and equilibrium-dependent models. As in Fig. 9, the difference was taken $t_{\text{model}} - t_{\text{actual}}$, so that positive values (towards red) indicate late predictions while negative values (towards blue) indicate early predictions. Inset partial histograms show the number of forecasts (out of 5000) predicting the first flow reversal within 3 minutes of the truth (green) and predicting that a flow reversal will never happen (red); the axes have the same scale for comparison, and the bar colors correspond to the dot colors. (B) Applying lobe-dependent correction (as opposed to (A) dynamically uninformed correction) increases the number of initial states for which the first flow reversal is predicted accurately, visible as a larger number of green dots (bigger green bar). Also, although the region of dark red (initial states attracted to convective equilibrium) in the left lobe has decreased in size, the one in the right lobe has inflated, again as compared to the top left. (bottom left) Equilibrium-dependent correction shrinks the left-lobe red region without inflating the one in the right lobe. However, it maintains the region of initial states for which flow reversal predictions are slightly early (light blue), which was reduced by the lobe-dependent correction. (bottom right) Applying the simultaneously lobe and equilibrium dependent correction, in which there are four different correction regions, the forecast model demonstrates the smallest region of initial states whose qualitative dynamics are different from the CFD simulated thermosyphon.

V. CONCLUDING DISCUSSION

It is apparent from the results of this work that the empirical correction technique tested here is successful in reducing forecast error in the low-dimensional setting. In both the perfect model scenario and the mock-operational experiment, improved error statistics and prolonged usefulness of forecasts were demonstrated by the corrected models. Furthermore, the empirical cor-

rection procedure was shown to provide greater improvement in average forecast accuracy than fine-tuning of parameters. However, this positive result comes at the cost of altering some important dynamical characteristics of the model. Details and implications of these results are discussed in the next two sections. In the third section, we make an observation about the relative impact of state-independent and state-dependent correction, in the present context, to conclude the paper.

A. Empirical correction vs. parameter tuning

In the perfect model scenario, empirical correction of models with the greatest parameter error provided forecasts with smaller short-term average error, and that were useful for longer, than those made by the uncorrected models with least parameter error. Similarly, empirically correcting the EM model with 10% error in every parameter (measured from tuned values) produced a forecast model that outperforms the tuned, but uncorrected model in predicting the CFD simulated thermosyphon. In each experiment, superior performance is evident in terms of anomaly correlation and average forecast error, measured against 3DVar analysis, and is verified against observed mass flow-rate of the CFD simulation in the toy climate experiment.

We do not present these results as evidence that empirical correction should replace parameter tuning, nor even that the former is better than the latter in any well-defined sense. However, the results do suggest that empirical correction could be a viable complement to the tuning of model parameters. Particularly as degrees of freedom become large, for example $N \approx 10^{10}$ in some currently operational numerical weather models, the computational cost of parameter tuning is very large in comparison to that of empirical correction, when appropriately modified for such models (see [9, 13]). An example of such a modification is to compute the first $m \ll N$ principal components (PCs) of \mathbf{L} by singular value decomposition (SVD), where m is determined so that a certain percentage of the state covariance is explained by these first m PCs. A combined tuning/correction approach could reduce the number of model integrations necessary in the parameter tuning process without sacrificing model accuracy. However, the dynamical ramifications of this strategy must be considered, as we explain in the next section.

B. Error statistics vs. system dynamics

The reduction in average forecast error provided by empirical correction belies fundamental dynamical disturbances born out of the correction procedure. Stabilization of equilibrium solutions, which results in large regions of initial states for which the corrected model forecasts behavior in opposition to that of the true system, follows from the dynamical modification imposed by empirical correction. In addition, the symmetry of the model system is broken by empirical correction. Though these costs can be mitigated somewhat by hard-wiring system-specific dynamical cues into the correction procedure, they can not be eradicated without more fundamental alterations of the technique, e.g., forcing the bias term \mathbf{b} and Leith operator \mathbf{L} to preserve system symmetry. In operational practice, empirical correction is known to introduce imbalances, e.g. violating geostrophy, necessitating some mechanism for smoothing the

flow into a physically viable region of state space.

In fact, it may be impossible to avoid all dynamical inaccuracies resulting from empirical correction, and even if theoretically possible, it would likely be impractical to do so in any operational setting. In considering the application of the technique in operational settings, then, it must be determined if the effects of misrepresented dynamics can be reduced to a tolerable level on a case-by-case basis. In the ideal situation, regions of state-space that would be dynamically misrepresented under empirical correction could be reduced, by minor modifications to the correction procedure, to encompass only unrealistic or unlikely physical states. In any case, the technique presented in this study should not be applied without such considerations.

C. Bias correction vs. Leith operator

State-independent error correction by itself produces almost no improvement in any of the forecast models in this study (not shown). This is in contrast to what has been observed in operational weather and climate model studies, where state-independent bias correction typically outperforms state-dependent correction in reduction of forecast errors [20]. The inaccuracies of ad-hoc forcings included in such models to compensate for external and/or irresolvable phenomena (e.g. solar and cloud forcings, respectively) are likely responsible for a large component of the bias. In light of the lack, or at least minimal nature of such external and sub-gridscale influences in the toy models considered here, the ineffectiveness of bias correction is logically consistent with this explanation.

The state-dependent Leith operator is entirely responsible for the success of the corrected models in this study. In the perfect model scenario this makes sense because the difference between the forecast models and the “truth” model are inherently multiplicative, i.e. the parameters are coefficients weighting the interaction between state-variable values and thus resulting errors *must* depend on state. For the EM model of the CFD system, it seems that errors resulting from the low dimensionality of the forecast model may also be multiplicative in nature. If this is the case, state-dependent correction may reduce error patterns in operational models that result from reduced dimensionality, e.g. coarse resolution. The correction will not likely compensate for processes that are irresolvable due to coarse resolution, but rather may reduce the propagation of error resulting from the omission of such phenomena. This hypothesis is consistent with demonstrated improvement of local behavior in state-dependent corrected atmospheric models with $N \approx 10^5$ degrees of freedom [9].

In previous studies of state-dependent correction in models that are much more realistic than those considered here, resulting error reduction has been minuscule in comparison to what is achieved by bias correction. How-

ever, this is not cause to reject the usefulness of parameterizing state-dependent error. Though globally averaged error reduction may not be significant, improvement in the local behavior of models can have a large impact on forecast uncertainty, particularly in an ensemble strategy where state-dependent correction can increase the spread in previously unsampled state-space directions.

ACKNOWLEDGEMENTS

Simulations and experiments were performed on the Vermont Advanced Computing Center (VACC) 1400-processor cluster, an IBM e1350 High Performance Computing system. The authors would like to thank Ross Lieb-Lappen for assistance in preparing this manuscript. This study was supported by NSF grant #DMS0940271, the Mathematics and Climate Research Network, and a National Aeronautics and Space Administration (NASA) EPSCoR grant.

-
- [1] J. Anderson, *Monthly Weather Review* **129**, 2884 (2001).
 - [2] J. Whitaker and T. Hamill, *Monthly Weather Review* **130**, 1913 (2002).
 - [3] E. Ott, B. R. Hunt, I. Szunyogh, A. V. Zimin, E. J. Kostelich, M. Corazza, E. Kalnay, D. J. Patil, and J. A. Yorke, *Tellus* **56**, 415 (2004).
 - [4] T. Hamill and C. Snyder, *Monthly Weather Review* **128**, 2905 (2000).
 - [5] P. Houtekamer and H. Mitchell, *Monthly Weather Review* **129**, 796 (2001).
 - [6] E. Kalnay, *Atmospheric Modeling, Data Assimilation and Predictability* (Cambridge University Press, 2003).
 - [7] C. E. Leith, *Annual Review of Fluid Mechanics* **10**, 107 (1978).
 - [8] T. DelSole and A. Y. Hou, *Monthly Weather Review* **127**, 2533 (1999).
 - [9] C. M. Danforth, E. Kalnay, and T. Miyoshi, *Monthly Weather Review* **135**, 281 (2007).
 - [10] E. N. Lorenz, *Journal of the Atmospheric Sciences* **20**, 130 (1963).
 - [11] E. N. Lorenz, *Tellus* **17**, 321 (1965).
 - [12] C. M. Danforth and E. Kalnay, *Geophysical Research Letters* **35**, L24805 (2008).
 - [13] C. M. Danforth and E. Kalnay, *Journal of the Atmospheric Sciences* **65**, 1467 (2008).
 - [14] P. Ehrhard and U. Müller, *Journal of Fluid Mechanics* **217**, 487 (1990).
 - [15] K. D. Harris, E.-H. Ridouane, D. L. Hitt, and C. M. Danforth, In preparation (2011).
 - [16] E. Evans, N. Bhatti, J. Kinney, L. Pann, M. Peña, S.-C. Yang, E. Kalnay, and J. Hansen, *Bulletin of the American Meteorological Society* **85**, 520 (2004).
 - [17] S.-C. Yang, D. Baker, H. Li, K. Cordes, M. Huff, G. Nagpal, E. Okereke, J. Villafañe, E. Kalnay, and G. S. Duane, *Journal of the Atmospheric Sciences* **63**, 2340 (2006).
 - [18] R. N. Miller, M. Ghil, and F. Gauthiez, *Journal of the Atmospheric Sciences* **51**, 1037 (1994).
 - [19] E. H. Ridouane, C. M. Danforth, and D. L. Hitt, *International Journal of Heat and Mass Transfer* **53**, 76 (2010), ISSN 0017-9310, URL <http://www.sciencedirect.com/science/article/B6V3H-4XKF9BN-2/2/08a02c9fae7bf2fbdae3390290fb129b>.
 - [20] X. Yang, T. DelSole, and H. L. Pan, *Monthly Weather Review* **136**, 5224 (2008).